PRODUCT NAME

Compaq TruCluster™ Server V5.1A

OVERVIEW

Description

Compaq TruCluster™ Server Version 5.1A for Compaq Tru64™ UNIX Version 5.1A, the latest in the new generation in the TruCluster product line from Compaq, provides highly available and scalable solutions for mission-critical computing environments. TruCluster Server delivers powerful but easy-to-use UNIX clustering capabilities.

By combining the advantages of symmetric multiprocessing (SMP), distributed computing, and fault resilience, a cluster running TruCluster Server offers high availability while providing scalability beyond the limits of a single system. On a single-system server, a hardware or software failure can severely disrupt a client's access to critical services. In a TruCluster Server cluster, a hardware or software failure on one member system results in the other members providing these services to clients.

By extending single-system management capabilities to clusters, TruCluster Server reduces the effort and complexity of cluster administration. It provides a clusterwide namespace for files and directories, including a single root file system that all cluster members share. A common cluster address (cluster alias) for the Internet protocol suite (TCP/IP) makes the cluster appear as a single system to its network clients while load balancing client connections across member systems.

A single system image allows a cluster to be managed more easily than distributed systems. TruCluster Server cluster members share a single root file system and common system configuration files. Therefore, most management tasks need to be done only once for the entire cluster rather than repeatedly for each cluster member. The cluster can be managed either locally from any of its members or remotely using Tru64 UNIX Web-based management tools. Tru64 UNIX and TruCluster Server software, and applications, are installed only once. Most network applications, such as the Apache Web server, need to be configured only once in the cluster and can be managed more easily in a cluster than on distributed systems.

A choice of graphical, Web-based, or command-line user interfaces makes management tasks easier for the administrator, flexible for those with large configurations, and streamlined for expert users.

TruCluster Server facilitates deployment of services that remain highly available even though they have no embedded knowledge they are running in a cluster. Applications can access their disk data from any cluster member. TruCluster Server also provides the support for components of distributed applications to run in parallel, providing high availability while taking advantage of cluster-specific synchronization mechanisms and performance optimizations.

TruCluster Server allows the processing components of an application to concurrently access raw devices or files, regardless of where the storage is located in the cluster. Member-private storage and clusterwide shared storage are equally accessible to all cluster members. Using either standard UNIX file locks or the distributed lock manager (DLM), an application can synchronize clusterwide access to shared resources, maintaining data integrity.

TruCluster Server is an efficient and reliable platform for providing services to networked clients. To a client, the cluster appears to be a powerful single-server system; a client is impacted minimally, if at all, by hardware and software failures in the cluster.

TruCluster Server simplifies the mechanisms of making applications highly available. A cluster application availability (CAA) facility records the dependencies of, and transparently monitors the state of, registered applications. If a hardware or software failure prevents a system from running a service, the failover mechanism automatically relocates the service to a viable system in the cluster, which maintains the availability of applications and data. Administrators can manually relocate applications for load balancing or hardware maintenance.

TCP-based and UDP-based applications can also take advantage of the cluster alias subsystem. These applications, depending on their specific characteristics, can run on a single cluster member or simultaneously on multiple members. The cluster alias subsystem routes client requests to any member participating in that cluster alias. During normal operations, client connections are dynamically distributed among multiple service instances according to administrator-provided metrics.

COMPAQ

OVERVIEW *(continued)*

Description *(continued)*

TruCluster Server supports a variety of hardware configurations that are cost-effective and meet performance needs and availability requirements. Hardware configurations can include different types of systems and storage units, and can be set up to allow easy maintenance. In addition, administrators can set up hardware configurations that allow the addition of a system or storage unit without shutting down the cluster.

For the fastest communication with the lowest latency, use the PCI-based Memory Channel cluster interconnect for communication between cluster members. TruCluster Server Version 5.1A also supports the use of 100 Mbps Ethernet hardware as a private LAN cluster interconnect. All cluster members must all use Memory Channel or all must use 100 Mbps LAN, but these interconnects cannot be mixed for cluster communication in the same cluster. Using multiple shared buses and redundant Memory Channel or LAN interconnect hardware promotes no-single-point-of-failure (NSPOF) characteristics for mission-critical applications.

FEATURES –TRUCLUSTER SERVER V5.1A

A TruCluster Server cluster acts as a single virtual system, even though it is made up of multiple systems. Cluster members can share resources, data, and clusterwide file systems under a single security and management domain, yet they can boot or shut down independently without disrupting the cluster.

Cluster File System

The Cluster File System (CFS) makes all files, including the root (/), /usr, and /var file systems, visible to and accessible by all cluster members. It does not matter whether a file is stored on a device connected to all cluster members or on one that is private to a single member. Each file system is served by a single-cluster member; other members access that file system as CFS clients with significant optimizations for shared access. CFS preserves full X/Open and POSIX semantics for file system access and maintains cache coherency across cluster members. For instance, an application can use standard UNIX file locks to synchronize access to shared files.

CFS supports the Advanced File System (AdvFS) for both read and write access and supports AdvFS with BSD-type user and group quotas. NFS client and NFS server are supported for both read and write access, and NFS services are accessible to clients through cluster aliases in addition to the default cluster alias. The UNIX File System (UFS) is supported for read and write access from the CFS server or for read-only access on client members. The Memory File System (MFS) is supported for read and write access on the CFS server. The CD-ROM File System (CDFS) and Digital Video Disc File System (DVDFS) are supported for read access only. For higher performance, applications can use direct I/O through the file system to bypass the buffer cache.

Device Request Dispatcher

The device request dispatcher supports clusterwide access to character and block disk devices, and to tape and tape changer devices. All local and remote cluster disk and tape I/O passes through the device request dispatcher. A member does not need a direct connection to a disk, or tape, or tape changer device to access data on that device. This permits great flexibility in selecting a hardware configuration that is both economical and useful.

Cluster Alias

A cluster alias is an IP address that makes the cluster look like a single system to clients and other hosts on the network. Cluster aliases free clients from having to connect to specific members for services. If the member providing the service goes down, a client reconnects to another member elected by the cluster alias to provide the service. With applications that run concurrently on multiple members, scaling is achieved by permitting multiple clients to connect to instances of the service on multiple cluster members, each using a cluster alias to address the service.

The cluster alias subsystem provides an optional virtual MAC (vMAC) address that can be associated with each cluster alias IP address. When configured, the same MAC address is used in all Address Resolution Protocol (ARP) responses for the cluster alias address, independent of which cluster node is responding to cluster alias ARP requests. This permits faster failover when a new node assumes responsibility for responding to cluster alias ARP requests.

Cluster Application Availability Facility

The cluster application availability (CAA) facility delivers the ability to deploy highly available single instance applications in a cluster by providing resource monitoring and application relocation, failover, and restart capabilities. CAA is used to define which members can run a service, the criteria under which to relocate a service, and the location of an application-specific action script. Monitored resources include network adapters, tape devices, media changers, and applications. CAA allows services to manage and monitor resources by using entry points within their action scripts. Applications do not need to be modified in any way to utilize CAA.

FEATURES –TRUCLUSTER SERVER V5.1A *(continued)*

Rolling Upgrade

TruCluster Server allows rolling upgrade from the previous version of the base operating system and the TruCluster software to the next subsequent release of the base operating system and TruCluster software. It also allows the rolling of patches into the cluster. Updating the operating system and cluster software does not require a shutdown of the entire cluster. A utility is provided to roll the cluster in a controlled and orderly fashion. The upgrade procedure allows the monitoring of the status of the upgrade while it is in progress. Clients accessing services are not aware that a rolling upgrade is in progress.

See the *Cluster Installation* manual for recommended and supported paths to upgrade or roll to the latest version of TruCluster Server.

Cluster Management

The SysMan system management utilities provide a graphical view of the cluster configuration, and can be used to determine the current state of availability and connectivity in the cluster. The administrator can invoke management tools from SysMan, allowing the cluster to be managed locally or remotely.

Clusterwide signaling allows applications to send UNIX signals to processes operating on other members.

Performance Management

The performance management capability of Tru64 UNIX has been modified from one large performance management tool (*pmgr*) to several smaller and more versatile tools. The performance management tool suite consists of: `Collect, Collgui`, and the two SNMP agents pmgrd and clu_mibs.

`Collect` is a tool that collects operating system and process data under Compaq Tru64 UNIX Versions 4.*x* and 5.*x*. Any subset of the 'subsystems' (Process, Memory, Disk, LSM Volumes, Network, CPU, Filesystems), and Header can be defined for which data is to be collected. `Collect` is designed for high reliability and low system-resource overhead. Accompanying `collect` are two highly integrated tools, `collgui`, a graphical front-end, and `cfilt`, which allows completely arbitrary extraction of data from the output of `collect` to standard output. `collgui` is a laborsaving tool that allows a user to quickly analyze collect data. The Performance Manager metrics server (*pmgrd*) is a UNIX daemon process that provides general UNIX performance metrics on request. The *pmgrd* metrics server supports the extensible Simple Network Management Protocol (SNMP) agent mechanism (eSNMP).

Cluster MIB

TruCluster Server supports the Compaq Common Cluster MIB. The Compaq Insight Manager uses this Cluster MIB to discover cluster member relationships, and to provide a coherent view of clustered systems across supported platforms.

Highly Available NFS Server

When configured as an NFS server, a TruCluster Server cluster can provide highly available access to the file systems it exports. There are no special cluster management operations required to configure the cluster as a highly available NFS server. In the event of a system failure, another cluster member will become the NFS server for the file system, transparent to external NFS clients. NFS file locking is supported, as are both NFS V2 and V3 with UDP and TCP.

TruCluster Server allows NFS file systems to be served from the cluster through both the default cluster alias and alternate aliases. Alternate cluster aliases can be defined to limit NFS server activity to those members that are actually connected to the storage that contains the exported file systems. NFS clients can use this alternate alias when they mount the file systems served by the cluster.

Fast File System Recovery

The Advanced File System (AdvFS) log-based file system provides higher availability and greater flexibility than traditional UNIX file systems. AdvFS journaling protects file system integrity. TruCluster Server supports AdvFS for both read and write access.

An optional, separately licensed product, the Advanced File System Utilities, performs online file system management functions. See the OPTIONAL SOFTWARE section of this document for more information on the AdvFS utilities.

Increased Data Integrity

Tru64 UNIX Logical Storage Manager (LSM) is a cluster-integrated, host-based solution to data storage management. In a TruCluster Server cluster, LSM operations continue despite the loss of cluster members, as long as the cluster itself continues operation and a physical path to the storage is available. LSM disk groups can be used simultaneously by all cluster members and the LSM configuration can be managed from any cluster member.

Basic LSM functionality, including disk spanning and concatenation, is provided with the Tru64 UNIX operating system. Extended functions, such as striping (RAID 0), mirroring (RAID 1), and online management, are available with a separate license. Mirroring of LSM is RAID Advisory Board (RAB) certified for RAID Levels 0 and 1.

LSM is supported for use in a TruCluster Server cluster and will support any volume in a cluster, including swap and cluster root and excluding the quorum disk and member boot disks. Hardware mirroring is supported for all volumes in a cluster without exception.

LSM RAID 5 volumes are not supported in clusters. See the OPTIONAL SOFTWARE section of this document for more information on LSM.

FEATURES –TRUCLUSTER SERVER V5.1A *(continued)*

Global Error Logger & Event Manager

TruCluster Server can log messages about events that occur in the TruCluster environment to one or more systems. Cluster administrators can also receive notification through electronic mail when critical problems occur.

Cluster Storage I/O Failover

TruCluster Server provides two levels of protection in the event of storage interconnect failure. When configured with redundant storage adapters, the storage interconnect will be highly available. Should one interconnect fail, traffic will transparently fail over to the surviving adapter. When a member system is connected to shared storage with a single storage interconnect and it fails, transactions are transparently performed via the cluster interconnect to another cluster member with a working storage interconnect.

Cluster Client Network Failover

TruCluster Server supports highly available client network interfaces via the Tru64 UNIX redundant array of independent network adapters (NetRAIN) feature.

Cluster Interconnect Failover

TruCluster Server allows the elimination of the cluster interconnect as a single point of failure by supporting redundant cluster interconnect hardware. You can configure dual-rail Memory Channel, allowing the cluster to survive the failure of a single rail. For LAN interconnect, two or more network adapters on each member are configured as a NetRAIN virtual interface. When properly configured across two or more switches, the cluster will survive any LAN component failure. This not only guards against rare network hardware failures, but also facilitates the upgrade and maintenance of the network without disrupting the cluster.

Support for Parallelized Database Applications

TruCluster Server provides the software infrastructure to support parallelized database applications, such as Oracle Parallel Server (OPS) and Informix Extended Parallel Server (XPS) to achieve high performance and high availability. OPS and XPS are offered and supported separately by Oracle Corporation and Informix Software, Inc., respectively.

Distributed Lock Manager

The distributed lock manager (DLM) synchronizes access to resources that are shared among cooperating processes throughout the cluster. DLM provides a software library with an expansive set of lock modes that applications use to implement complex resource-sharing policies. DLM provides services to notify a process owning a resource that it is blocking another process requesting the resource. An application can also use DLM routines to efficiently coordinate the application's activities within the cluster.

Support for Memory Channel API

TruCluster Server provides a special application programming interface (API) library for high-performance data delivery over Memory Channel by giving access to Memory Channel data transfer and locking functions. This Memory Channel API library enables highly optimized applications that require high-performance data delivery over the Memory Channel interconnect. This library is supported solely for use with Memory Channel.

High performance within the cluster is achieved by providing user applications with direct access to the capabilities of the Memory Channel. For example, a single store instruction on the sending host is sufficient for the data to become available for reading in the memory of another host.

The Memory Channel API library allows a programmer to create and control access to regions of the clusterwide address space by specifying UNIX-style protections. Access to shared data can be synchronized using Memory Channel spin locks for clusterwide locking.

The Memory Channel API library facilitates highly optimized implementations of Parallel Virtual Machine (PVM), Message Passing Interface (MPI), and High Performance Fortran (HPF), providing seamless scalability from SMP systems to clusters of SMP machines. This provides the programmer with comprehensive access to the current and emerging de facto standard software development tools for parallel applications while supporting portability of existing applications without source code changes.

Note to users of the Memory Channel API V1.6 product on Tru64 UNIX Version 4.0*:

On Tru64 UNIX V5.*, the TruCluster Memory Channel Software V1.6 product is bundled as a feature of TruCluster Server V5.*. To run the Memory Channel API library on Tru64 UNIX version 5.1A, you must install a TruCluster Server license to configure a valid TruCluster Server cluster.

FEATURES –TRUCLUSTER SERVER V5.1A *(continued)*

Connection Manager

The connection manager is a distributed kernel component that ensures that cluster members communicate with each other and enforces the rules of cluster membership. The connection manager forms a cluster, and adds and removes cluster members. It tracks whether members in a cluster are active and maintains a cluster membership list that is consistent on all cluster members.

Support for Fibre Channel Solutions

TruCluster Server supports the use of switched Fibre Channel storage and Fibre Channel Arbitrated Loop. Fibre Channel provides the following benefits over parallel SCSI storage:

- Greater performance
- Greater scalability
- More reliability, serviceability, and availability

Fibre Channel can be used for clusterwide shared storage, cluster file systems, swap partitions, and boot disks. Compared to parallel SCSI storage, Fibre Channel is easier to configure and its long distance permits greater flexibility in configurations.

Compared with a switched Fibre Channel topology, Arbitrated Loop offers a lower cost solution by trading off bandwidth, and therefore some performance. Arbitrated Loop is supported for two-member configurations only.

For more information on supported TruCluster Server configurations and specific cabling restrictions using Fibre Channel, see the *Cluster Hardware Configuration* manual at the following URL: http://tru64unix.compaq.com/docs/pub_page/cluster_list.html

Enhanced Security with Distributed Authentication

TruCluster Server supports the Enhanced Security option on all cluster members. This includes support for features for enhanced login checks and password management. Audit and access control list (ACL) support can also be enabled independently of the Enhanced Security option on cluster members.

CONFIGURATION

Software Requirements

TruCluster Server Version 5.1A requires the Tru64 UNIX Version 5.1A operating system. The Tru64 UNIX operating system is a separately licensed product. See the Tru64 UNIX operating system QuickSpec for more information.

Software Configuration Requirements

- When configuring TruCluster Server software an additional 64 MB of memory is required on each member system.
- Each system requires a disk for a member boot disk and the cluster requires a minimum of one disk for clusterwide root, /usr, and /var file systems. A quorum disk is optional. See the TruCluster Server *Cluster Hardware Configuration* manual for details. Free disk space required for use (permanent): 52 MB to load TruCluster Server software onto a Tru64 UNIX system disk.

These requirements refer to the disk space required on the system disk. The sizes are approximate; actual sizes may vary depending on the system environment, configuration, and software options.

ORDERING INFORMATION

TruCluster Server V5.1A licenses include:

Systems	TruCluster Plus packages*	TruCluster Server License	TruCluster Server Migration License**
AlphaServer™ 800, 1000A, 1200, DS10, DS10L, DS20, DS20E	QP-6R9AC-AA	QL-6BRAC-AA	QL-6J1AC-AA
AlphaServer ES40	QP-6R9AE-AA	QL-6BRAE-AA	QL-6J1AE-AA
AlphaServer 2000, 2100, 2100A, 4000, 4100, GS60E, GS80	QP-6R9AG-AA	QL-6BRAG-AA	QL-6J1AG-AA
AlphaServer 8200, 8400, GS60, GS140, GS160, GS320	QP-6R9AQ-AA	QL-6BRAQ-AA	QL-6J1AQ-AA

Software Documentation: QA-6BRAA-GZ

*TruCluster Plus Software packages include licenses for: TruCluster Server, Logical Storage Manager, and AdvFS Utilities.

**If you currently have TruCluster Available Server or TruCluster Production Server and want to convert to TruCluster Server, use the QL-6J1A*-AA migration license.

Software Licensing

TruCluster Server is licensed as a standard Compaq software layered product and is provided only under a license. Each system in the TruCluster Server environment requires separate Tru64 UNIX and TruCluster Server licenses. Hard partitions of GS80, GS160 and GS320 can be clustered together either across separate systems or within systems and only one TruCluster license is required per system. For more information about the Compaq licensing terms and policies, contact your local Compaq representative or reseller.

This product supports the Tru64 UNIX License Management Facility (LMF). License units for the TruCluster Server product are allocated on an unlimited-system-use basis.

For more information on the License Management Facility, see the Tru64 UNIX Operating System QuickSpec or the Tru64 UNIX operating system documentation.

Distribution Media

TruCluster Server is a separately licensed product and is distributed on the Tru64 UNIX Associated Products Volume 2 CD-ROM. The TruCluster Documentation is available on-line at http://www.tru64unix.compaq.com/docs/.

Software Product Services

A variety of service options are available from Compaq. For more information, contact your local Compaq service representative.

OPTIONAL SOFTWARE

Compaq Advanced Server for UNIX (ASU)

Compaq Advanced Server for UNIX (ASU) provides Windows networking services, such as file sharing, print sharing, and security for Tru64 UNIX. In addition to basic file and print services, ASU provides full Windows domain controller support, support for enterprise-wide trust relationships, and support for Windows security — including file permissions and Windows local and global groups.

Additionally, you can manage users, file shares, and printers using native Windows administrative tools. When combined with TruCluster Server software, ASU provides highly available and highly scalable file shares, print shares, and even Primary Domain Controller resources to Windows clients. For more information on ASU, visit the Compaq Advanced Server for UNIX Web site at:

www.tru64unix.compaq.com/products/adv_server

Legato NetWorker

See the *NetWorker Read This First* letter for information on evaluating or purchasing a version of NetWorker that supports TruCluster Server.

Compaq SANworks™ Data Replication Manager

Compaq SANworks Data Replication Manager (DRM) is controller-based data replication software for disaster tolerance and data movement solutions. It works with the new StorageWorks™ Fibre Channel MA8000 /EMA12000 Storage Solutions. Multiple clusters or standalone systems can be interconnected using DRM. The RAID Array 8000 (RA8000) and Enterprise Storage Array 12000 (ESA12000) are also supported.

For more information, see the SANworks Data Replication Manager QuickSpec.

Advanced File System (AdvFS) Utilities

The Advanced File System (AdvFS) og-based file system provides flexibility, compatibility, high availability, and high performance for files and filesets, up to 16 terabytes (TB). Administrators can add, remove, reconfigure, tune, and defragment files — and back up storage — without unmounting the file system or halting the operating system. By supporting multivolume file systems, AdvFS enables file-level striping to improve file transfer rates, and integrates with the functionality provided by the Logical Storage Manager (LSM).

A graphical user interface simplifies management tasks to dynamically resize file systems, load balance, undelete files, and clone files for hot backup.

The AdvFS Utilities is a separately licensed software product for Tru64 UNIX. See the AdvFS Utilities QuickSpec for more information.

Logical Storage Manager (LSM)

The Tru64 UNIX Logical Storage Manager (LSM) is an integrated, host-based solution to data storage management. Concatenation, striping, mirroring, hot-sparing, and a graphical user interface allow data storage management functions to be done online, without disrupting users or applications. LSM manages storage as a single entity in both cluster and single-node environments. LSM is a separately licensed software product for Tru64 UNIX. For more information, see the Logical Storage Manager QuickSpec.

***StorageWorks*™ Software**

The StorageWorks Software package includes the licenses for Tru64 UNIX Logical Storage Manager and the Advanced File System Utilities. The part number for the StorageWorks software package is QB-5RXA*-AA.

TruCluster Plus Software

The TruCluster Plus Software package includes the licenses for TruCluster Server, Tru64 UNIX Logical Storage Manager, and the Advanced File System Utilities. The part number for the TruCluster Plus software package is QP-6R9A*-AA.

OPTIONAL SOFTWARE *(continued)*

Service Tools - WEBES

WEB-Based Enterprise Service (WEBES) tools integrate a high availability system fault management architecture, *Distributed Enterprise Service Tools Architecture* (DESTA), with Compaq's architecture for distributed, Web-Based System Management. The tool functionality contained in the WEBES kit includes the following: *Compaq Analyze* (symptom directed hardware diagnosis tool), *Compaq Crash Analysis Tool* [CCAT] (symptom directed operating system software diagnosis tool), and *Revision and Configuration Management* [RCM] tool (collects system configuration and revision data).

Compaq Analyze is a hardware diagnosis software tool that provides analysis for single errors or fault events at a rudimentary level, as well as multiple event and complex analysis. *Compaq Analyze* provides automatic notification and isolation of hardware components to quickly identify areas of the system that may be having problems. *Compaq Analyze* is the successor to *DECevent* and supports the newer EV6 based systems. Refer to the release notes for the products that are supported.

CCAT is a software application tool that helps service engineers and system managers to analyze operating system crashes. This tool collects data that describes system crashes and matches that data against a set of operating specific rules.

The RCM tool collects system configuration and revision data information. The data is stored in the RCM Server at Compaq Services and the server is then used to create detailed revision and configuration reports.

Oracle Parallel Server (OPS)

Oracle Parallel Server (OPS) technology is a relational database management system that capitalizes on the benefits of high availability, performance, and expandability made possible by Tru64 UNIX clusters. Oracle Parallel Server must be ordered separately through Oracle Corporation.

Informix Dynamic Server

Informix Dynamic Server with the Extended Parallel Option delivers high performance while supporting commercially available cluster systems used for data warehousing and decision support applications. Informix Dynamic Server with the Extended Parallel Option must be ordered through Informix Software, Inc.

SUPPORTED SYSTEMS AND CLUSTER INTERCONNECT HARDWARE REQUIREMENTS

TruCluster Server Version 5.1A supports the systems listed in the following table with up to eight systems in a configuration.
TruCluster Server Version 5.1A supports the KZPSA-BB, KZPBA-CB, KGPSA-BC, and KGPSA-CA as shared storage bus adapters, subject to the current maximum number of adapters and any other restrictions for a given system. TruCluster Server supports a maximum of 62 shared buses in any combination. Information on firmware release can be found at http://www.compaq.com/support/ or from the current Alpha™ systems firmware update CD-ROM.

Supported Systems

System	Shared Storage IO Adapter	Cluster Interconnect
AlphaServer 800	KZPSA, KZPBA, KGPSA	MC1, 1.5 & 2, 100 Mbps LAN
AlphaServer 1000A	KZPSA	MC1.5, 100 Mbps LAN
AlphaServer 1200	KZPSA, KZPBA, KGPSA	MC1, 1.5 & 2, 100 Mbps LAN
AlphaServer 2000	KZPSA	MC1, 1.5, 100 Mbps LAN
AlphaServer 2100	KZPSA	MC1, 1.5 , 100 Mbps LAN
AlphaServer 2100A	KZPSA	MC1, 1.5 , 100 Mbps LAN
AlphaServer 4000, 4100	KZPSA, KZPBA, KGPSA	MC1, 1.5 & 2, 100 Mbps LAN
AlphaServer 8200	KZPSA, KZPBA, KGPSA	MC1, 1.5 & 2, 100 Mbps LAN
AlphaServer 8400	KZPSA, KZPBA, KGPSA	MC1, 1.5 & 2, 100 Mbps LAN
AlphaServer DS10	KZPBA, KGPSA	MC2, 100 Mbps LAN
AlphaServer DS10L	KZPBA	100 Mbps LAN
AlphaServer DS20, DS20E	KZPBA, KGPSA	MC2, 100 Mbps LAN
AlphaServer ES40	KZPBA, KGPSA	MC2, 100 Mbps LAN
AlphaServer GS60, GS60E	KZPSA, KZPBA, KGPSA	MC1, 1.5 & 2, 100 Mbps LAN
AlphaServer GS80	KZPBA, KGPSA-CA	MC2, 100 Mbps LAN
AlphaServer GS140	KZPSA, KZPBA, KGPSA	MC1, 1.5 & 2, 100 Mbps LAN
AlphaServer GS160	KZPBA, KGPSA-CA	MC2, 100 Mbps LAN
AlphaServer GS320	KZPBA, KGPSA-CA	MC2, 100 Mbps LAN

Notes:

- Hard partitions of GS80, GS160, and GS320 can be clustered together either across separate systems or within systems, but each hard partition must have at least one cluster interconnect connection.
- TruCluster Server Version 5.1A requires all members to be connected with either all members using Memory Channel hardware or all members using a private 100 MB full-duplex LAN. Note that there are two variants of Memory Channel (MC):

SUPPORTED SYSTEMS AND CLUSTER INTERCONNECT HARDWARE REQUIREMENTS *(continued)*

Supported Memory Channel Hardware

MC 1 & 1.5	MC2	Description
CCMAA-AA or CCMAA-BA	CCMAB-AA	PCI Adapter
CCMHA-AA	CCMHB-AA	Hub
CCMLA-AA	CCMLB-AA	Line Card
N/A BC12N-10	BN39B-04	4-meter Cable 10-foot copper cable
	BN39B-10	10-meter Cable
	BN39B-01 (one meter)	Connects MC adapter to CCMFB optical converter
	CCMFB-AA	Fibre Optics Converter
	BN34R-10 (10 meter)	Fibre-optic cable: Connect one optical converter to another
	BN34R-31 (31 meter)	Fiber-optic cable: Connect one optical converter to another

Memory Channel configuration notes:

- At least one Memory Channel adapter must be installed in a PCI slot in each member system. One or more link cables is required to connect systems to each other or to a hub. A cluster environment with two nodes does not require a hub. A configuration of more than two members requires a Memory Channel hub.
- MC1and MC1.5 cannot be mixed on the same rail with MC2.

There are special rules about circumstances where Memory Channel 1 and Memory Channel 2 can be used together in the same cluster. The TruCluster Server *Cluster Hardware Configuration* manual provides information regarding supported Memory Channel configurations.

Supported LAN Interconnect Hardware

Use of dedicated 100 MB LAN is supported for use as a cluster interconnect. A LAN interconnect must be private to cluster members. As long as any packet that is transmitted by a cluster member's interconnect adapter can only be received by interconnect adapters of other members of the same cluster, the interconnect meets the privacy requirement.

A LAN interconnect can be a direct connection between two cluster members or can employ hubs or switches. In general, any Ethernet adapter, switch, or hub that works in a standard LAN at 100 Mbps should work within a LAN interconnect. (Fiber Distributed Data Interface (FDDI), ATM LAN Emulation (LANE), 10 Mbps Ethernet, and Gigabit Ethernet are not supported.)

Although hubs and switches are interchangeable in most LAN interconnect configurations, switches are recommended for performance and scalability. Most hubs run in half-duplex mode and do not detect network collisions, so their use in a LAN interconnect may limit cluster performance. Overall, using a switch, rather than a hub, provides greater scalability for clusters with three or more members.

Adapters and switch ports must be configured compatibly with respect to speed (100 Mbps) and operational mode (full-duplex). A maximum of three hops is allowed between cluster members, where a hop means passing from a system, switch, hub, or router, to another system, switch, hub, or router. That is, any combination of up to two hubs, switches, or routers is supported between two cluster members. You must not introduce unacceptable latencies by using, for example, a satellite uplink or a wide area network (WAN) in the path between two components of a LAN interconnect.

A fully redundant LAN interconnect configuration employs two or more Ethernet adapters in a NetRAIN set on each member, with redundant wiring to two or more switches interlinked by two crossover cables. These Ethernet switches must be capable of one of the follow mechanisms for managing traffic across parallel inter-switch links: link aggregation (also known as port trunking), resilient links, or per-port-enabled spanning tree algorithm.

STORAGE HARDWARE REQUIREMENTS

Supported Fibre Channel Hardware

TruCluster Server supports the following Fibre Channel hardware. For more information on the supported Fibre Channel solutions, see the TruCluster Server *Release Notes* and the *Cluster Hardware Configuration* manual.

Device	Description
KGPSA-BC	PCI-to-Fibre Channel host adapter
KGPSA-CA	PCI-to-Fibre Channel host adapter
DS-DSGGA-AA/AB	8/16-port Fibre Channel switch
DS-DSGGB-AA/AB	8/16-port Fibre Channel switch
DS-DSGGC-AA/AB	8/16-port FC switch
DS-SWXHB-07	7-port FC hub
HSG60	Array Controller
HSG80	Array Controller

HSG60 and HSG80 controllers may be contained in many cabinet configurations including MA6000, RA8000, MA8000, ESA12000, EMA12000, and EMA16000.

Supported SCSI Controllers

The systems are connected to shared SCSI buses using an adapter from the following list. The TruCluster Server *Cluster Hardware Configuration* manual and *Release Notes* provide more information regarding SCSI controller configuration, respectively:

- KZPSA — PCI to Fast Wide Differential SCSI-2 adapter
- KZPBA-CB — PCI to UltraSCSI wide differential adapter

Supported SCSI Signal Converters

TruCluster Server supports the following SCSI signal converters. The TruCluster Server Hardware Configuration manual provides information regarding supported SCSI signal converters configuration.

Signal Converter	Description
DWZZA-AA	Standalone unit, single-ended/narrow to differential/narrow
DWZZA-VA	SBB, single-ended/narrow to differential/narrow
DWZZB-AA	Standalone unit, single-ended/wide to differential/wide
DWZZB-VW	SBB, single-ended/wide to differential/wide
DS-DWZZH-03	UltraSCSI hub
DS-DWZZH-05	UltraSCSI hub

Supported SCSI Cables

TruCluster Server supports the following SCSI cables:

Device	Description
BN21W-0B	SCSI-2 Cable “Y”
BN21R or BN23G	SCSI-2 Cable “A”
BN21K, BN21L, or 328215-00X	SCSI-3 Cable “P”
BN21M	50-pin LD to 68-pin HD
BC06P or BC19J	50-pin LD Cable
BN38C, BN38D, or BN38E	VHDCI-to-HD68 cable
BN37A	Ultra VHDCI Cable
BN37B	VHDCI-to-HD68 cable
BN21M	50-pin LD-to-HD68 cable
199629-002 or 189636-002	50-pin HD-to-68-pin HD
146745-003 or 146776-003	50-pin HD-to-50-pin HD
189646-001 or 189646-002	68-pin HD
BN38E-0B	HD68-to-VHDCI technology adapter cable

HARDWARE REQUIREMENTS

Supported Terminators and Connectors

TruCluster Server supports the following terminators and connectors:

Device	Description
H879-AA or 330563-001	HD68 terminator
H885-AA	HD68 tri-link connector
H8861-AA	VHDCI tri-link connector
H8863-AA	VHDCI terminator
H8574-A	50-pin LD terminator
H8860-A	50-pin LD terminator
341102-001	50-pin HD terminator
152732-001	VHDCI 68-pin LVD terminator

Supported Disk Devices

Every SCSI and Fibre Channel storage disk currently sold by Compaq that appears in the supported options list for a supported AlphaServer are supported for use in a cluster on a shared bus.

Some legacy disk devices are not supported for use on a shared bus. TruCluster Server supports the following list of disk devices on shared storage. Any SCSI or Fibre Channel storage disk manufactured by Compaq later than the ones in the following list, and appearing in the supported options list for a supported AlphaServer, are supported.

- RZ26-VA Narrow
- RZ26L-VA Narrow
- RZ26L-VW Wide
- RZ26N-VA Narrow
- RZ26N-VW Wide
- RZ28-VA Narrow
- RZ28-VW Wide
- RZ28B-VA Narrow
- RZ28D-VA Narrow
- RZ28D-VW Wide
- RZ28L-VA Narrow
- RZ28L-VW Wide
- RZ28M-VA Narrow
- RZ28M-VW Wide
- RZ29-VA Narrow
- RZ29-VW Wide
- RZ29B-VA Narrow
- RZ29B-VW Wide
- RZ29L-VA Narrow
- RZ29L-VW Wide
- RZ40-VA Narrow
- RZ40-VW Wide
- RZ40L-VA Narrow
- RZ40L-VW Wide

HARDWARE REQUIREMENTS *(continued)*

Supported Tape Devices and Media Changers

TruCluster Server supports the configuration of specific tape devices on a shared SCSI bus and on Fibre Channel. These devices will function properly in a multi-initiator environment. On a SCSI bus, they will be disrupted by bus resets that occur during cluster membership change events. Backup software must be explicitly capable of handling and recovering from such events and must utilize the cluster application availability (CAA) facility to facilitate highly available backup.

The following tape devices are supported on shared storage:

- TZ88 Tape Drive
  - Tabletop: TZ88N-TA
  - SBB: TZ88N-VA
- TZ89 Tape Drive
  - Tabletop: DS-TZ89N-TA
  - SBB: DS-TZ89N-VW
- TZ885 DLT MiniLibrary
  - Tabletop: TZ885-NT
  - Rackmount: TZ885-NE
- TZ887 DLT MiniLibrary
  - Tabletop: TZ887-NT
  - Rackmount: TZ887-NE
- Compaq 20/40-GB DLT Tape Drive
  - Tabletop: 340744-B21
- Compaq 40/80-GB DLT Tape Drive
  - Tabletop: 146197-B22
- TL891 DLT MiniLibrary (2-5-2 Part Numbers)
  - Tabletop: DS-TL891-NT[1]
  - Rackmount: DS-TL891-NE/NG[1,2]
- TL891 DLT MiniLibrary (6-3 Part Numbers)
  - Tabletop: TL891 with 1 DLT 35/70 drive — 120875-B21[3]
    TL891 with 2 DLT 35/70 drives — 120875-B22
  - Rackmount: TL891 with 1 DLT 35/70 drive — 120876-B21[3]
    TL891 with 2 DLT 35/70 drives — 129876-B22
    MiniLibrary Expansion Unit — 120877-B21[4]
    MiniLibrary Data Unit — 128670-B21
- TL881 DLT MiniLibrary
  - Tabletop: TL881 with 1 DLT 20/40 drive — 128667-B21[5]
    TL881 with 2 DLT 20/40 drives — 128667-B22[5]
  - Rackmount: TL881 with 1 DLT 20/40 drive — 128669-B21
    TL881 with 2 DLT 20/40 drives — 128669-B22[4]
    MiniLibrary Expansion Unit — 120877-B21
    MiniLibrary Data Unit — 128670-B21

- TL893 Automated Tape Library — DS-TL893-BA
- TL894 Automated Tape Library — DS-TL894-BA
- TL895 DLT Automated Tape Library

Number of Drives	2-5-2 Part Number	6-3 Part Number
2	DS-TL895-H2	349350-B22
3	N/A	349350-B23
4	NA	349350-B24
5	DS-TL895-BA	349350-B25
6	N/A	349350-B26
7	N/A	349350-B27
1	DS-TL89X-UA[6]	349351-B21[7]

- TL896 DLT Automated Tape Library — DS-TL896-BA
- ESL9326D Enterprise Library

Number of Drives	6-3 Part Number
0	146205-B21[8]
6	146205-B23[8]
8	146205-B24[8]
10	

1. Can be expanded to two TZ89N-AV drives with part number DS-TL892-UA.
2. A DS-TL890-NE/NG MiniLibrary Expansion Unit can be connected to up to three DS-TL891-NE/NG drive units to manage the drives and cartridges in all connected units. A DS-TL800-AA pass-through mechanism is required for the second and third DS-TL891-NE/NG.
3. Can be expanded to two DLT 35/70 drives with part number 120878-B21.
4. The MiniLibrary Expansion Unit can be used to control the drives and cartridges of up to five drive and data units. A MiniLibrary Pass-Through Mechanism, part number 120880-B21, is needed for each additional unit beyond the first drive unit.
5. Can be expanded to two DLT 20/40 drives with part number 128671-B21.
6. An upgrade kit to add one DS-TZ89N-VA tape drive.
7. An upgrade kit to add one 35/70 DLT tape drive.
8. Can be upgraded by one 35/70 DLT tape drive with part number 146209-B21.

HARDWARE REQUIREMENTS *(continued)*

Supported Tape Devices and Media Changers *(continued)*

Notes: The TL881, TL891, TL893, TL894, TL895, TL896, Compaq 20/40-GB DLT Tape Drive, Compaq 40/80 GB DLT Tape Drive, and ESL9326D Enterprise Library are supported with both the KZPSA and KZPBA-CB adapters. All other tape device and media changer support is provided with KZPSA adapters only.

The TL891 Mini Library, TL895 Automated Tape Library, and the ESL9326D Enterprise Library are also supported on a Fibre Channel storage bus with the KGPSA-BC and KGPSA-CA adapter.

The Compaq Fibre Channel Tape Controller and Modular Data Router are Fibre Channel-to-SCSI bridges that allow a SCSI tape device to communicate with other devices on a Fibre Channel.

For a list of tape devices supported by Enterprise Backup Storage, see the Enterprise Backup Storage documentation.

Supported Storage Boxes

TruCluster Server supports the following storage boxes:

Storage Box	Description
BA350	Single ended, narrow
BA356	Single ended, wide
DS-BA356	Ultra SCSI, SBB shelf

Supported Parallel SCSI Array (RAID) Controllers

TruCluster Server supports the following parallel SCSI RAID controllers on shared storage buses. The TruCluster Server *Release Notes* provide information regarding supported Array (RAID) controllers firmware revisions.

RAID Controller
SWXRA-Z1 Array Controller (HSZ20)
RAID Array 3000 (HSZ22)
HSZ40-Bx Array Controller
HSZ40-Cx Array Controller
HSZ50-Ax Array Controller
HSZ70 Array Controller
HSZ80 Array Controller

Network Adapters

TruCluster Server supports for client networks any Ethernet, FDDI, ATM (LAN emulation mode only), or Gigabit Ethernet adapters that are supported by the version of Tru64 UNIX on which it is running.

Hardware Restrictions

TruCluster Server has the following hardware restrictions. The TruCluster Server *Cluster Hardware Configuration* manual provides additional information regarding hardware restrictions.

- Prestoserve NVRAM failover is not supported on shared disk devices.

Hardware Configuration Examples

The TruCluster Server *Cluster Hardware Configuration* manual provides hardware configuration examples. TruCluster Server supports up to eight member cluster configurations as follows:

- **Switched Fibre Channel**: Eight-member systems may be connected to common storage over Fibre Channel in a fabric (switch) configuration.
- **Parallel SCSI**: Only four of the member systems may be connected to any one SCSI bus. Multiple SCSI buses may be connected to different sets of members and the sets of members may overlap. Use of a DS-DWZZH-05 UltraSCSI hub with fair arbitration enabled is recommended when connecting four member systems to a common SCSI bus.